

The Reality Gap in Dependability Research

FT algorithms/protocols research
vs.
Enterprise computing

Matti Hiltunen
AT&T Labs - Research
hiltunen@att.com

Disclaimer: Opinions expressed reflect the view of the author only, not those of AT&T.

SLAs

6σ \$\$\$

99.999

Customer
satisfaction

Complex systems

Transient failures

Performance failures

Multiple administrative
domains

Recovery/

Monitoring Repair

Maintenance

F out of N failures

Impossibility results

Synchronous vs
asynchronous

Failure models

Independent failures

Byzantine failures

Safety

◇W etc

Active Replication

Why?

Are the enterprise computing problems too easy (not exciting enough) to be of interest for research community?

- Typically crash or timing failures
- Not much active replication
- Hard to generalize into a nice abstract problem
- Are they not appreciated enough (publishable)?

Or are they too hard?

- Hard to assume failure independence
- Often transient failures
- Very complex system structure

Or is enterprise computing simply running too far behind (5-10 years)?

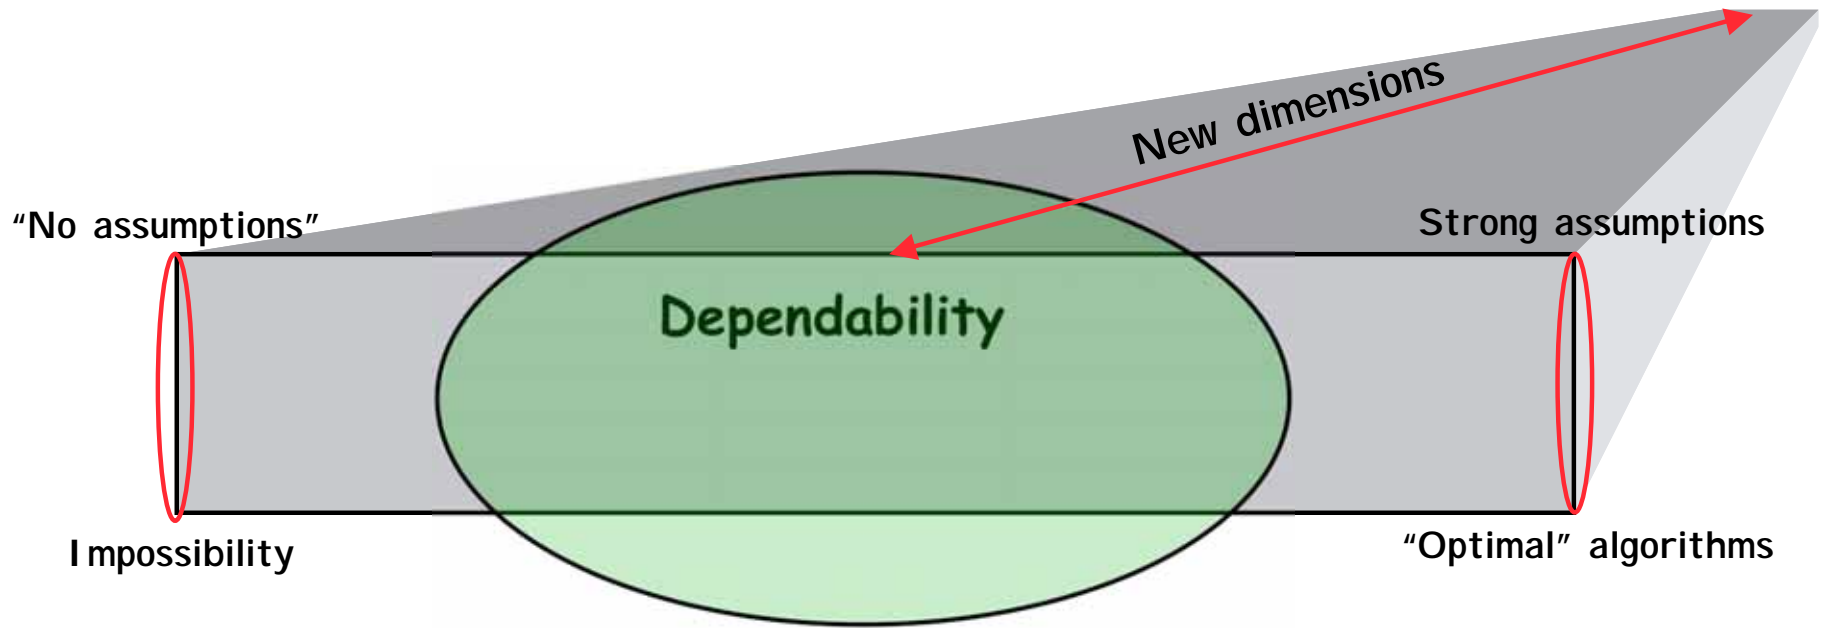
Or ...

Gaps

- Researchers like to solve nice abstract problems, but real systems are messy!
 - Multi-layer/tier systems
 - Network is more than a cloud
 - Potentially partially unknown topology
 - Multiple admin domains
 - Systems with COTS components
 - Shared computing resources (multiple apps, users)
 - Dependent failures
- Software failures/vulnerabilities are typically messy.
- Reliability and availability are probabilities
 - most research abstracts away probabilities



My view



Opportunities

- Low cost fault tolerance techniques (e.g., RAID)
 - Massive or even moderate replication not feasible (power, space, management, \$\$\$)
- Application-specific dependability techniques.
 - Potential to optimize; engineering challenges
- Diagnosis of transient and performance failures
- Automatic management (recovery, resource allocation)
- Consider different techniques (machine learning, control theory).
- Model-based algorithms can deal with probabilities
- More emphasis on “practical experience” and “engineering” papers.
 - Even negative results should be valuable.

